UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/575,071	04/07/2006	John A. Landis	10/777,000/TN333B	9251

7590 06/23/2010

Michael B Atlass
Unisys Corporation
Unisys Way
MS E8 114
Blue Bell, PA 19424-0001

EXAMINER	
CHEW, BRIAN	
ART UNIT	PAPER NUMBER
2195	

MAIL DATE	DELIVERY MODE
06/23/2010	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary	Application No.	Applicant(s)	
	10/575,071	LANDIS ET AL.	
	Examiner	Art Unit	
	BRIAN CHEW	2195	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *07 April 2006*.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-15* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-15* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☒ The specification is objected to by the Examiner.

10)☒ The drawing(s) filed on *7 April 2006* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)
2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date *4/7/2006*.
4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .
5) ☐ Notice of Informal Patent Application
6) ☐ Other: _____.

DETAILED ACTION

1. Claims 1-15 are presented for examination.

Specification

2. The disclosure is objected to because it contains an embedded hyperlink and/or other form of browser-executable code in paragraphs 41 and 51. Applicant is required to delete the embedded hyperlink and/or other form of browser-executable code. See MPEP § 608.01.

Claim Objections

3. Claim 13 is objected to because of the following informalities:

i. As per claim 13, Lines 3-4: There appears to be a typographical error at "de-queuing the partition relative physical addresses are converting the partition relative physical addresses". This shall be amended to read -- de-queuing the partition relative physical addresses and converting the partition relative physical addresses --. Appropriate correction is required.

Claim Rejections - 35 USC § 112

The following is a quotation of the second paragraph of 35 U.S.C. 112:

> The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

4. Claims 1-15 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. The language in the following claims is not clearly understood.

i. As per claim 1: Lines 1-2: It is uncertain what is meant by "A virtualization system for a host computer having at least one host processor and system resources including physical I/O hardware and memory" (i.e. Is the host computer part of the virtualization system?). Lines 7-8: It is uncertain what is meant by "maps said physical I/O hardware to endpoints of an I/O channel server" (i.e. What are endpoints of an I/O channel server? Do these correspond to the guest partitions or how the guest partitions access the I/O?). Lines 11-12: It is uncertain what is meant by "a resource database for use in managing use of said at least one host processor and said system resources" (i.e. How is the resource database used? Where is it located? In each guest partition or in the monitor?). Lines 14-15: It is uncertain what is meant by "guest applications in said at least one guest partition within memory space specified in said resource database" (i.e. Does this relate to less privileged user memory?).

ii. As per claim 2: Lines 3-4: It is uncertain what is meant by "a monitor associated with said at least one guest partition" (i.e. Does this refer to the at least one monitor of claim 1?).

iii. As per claim 3, it is dependent on claim 1 but does not overcome the deficiencies of claim 1; therefore, it is rejected for the same reasons.

iv. As per claim 4: Lines 1-2: It is uncertain what is meant by "an I/O monitor associated with said I/O partition" (i.e. Does this refer to the at least one monitor of claim 1?). Line 3: It is uncertain what is meant by "system call interface converting and validating client partition relative addresses" (i.e. Is this claiming the system call interface or the steps of converting and validating addresses? Does client partition refer to guest partition?). Lines 3-4: It is uncertain what is meant by "client memory channel drivers" (i.e. Does this relate to the memory channel comprising memory shared between the I/O and guest partition?). Lines 4-5: It is uncertain what is meant by "the client partition" (i.e. Does this refer to the guest partition?).

v. As per claim 5: Lines 1-2: It is uncertain what is meant by "a server of said I/O partition" (i.e. Does this refer to the I/O channel server of claim 1?). Lines 4: It is uncertain what is meant by "the I/O monitor" (i.e. Does this refer to the at least one monitor of claim 1?). Lines 5-6: It is uncertain what is meant by "whereby data may be exchanged with hardware I/O adapters connected between said I/O monitor and said common I/O physical hardware" (i.e. Is data actually exchanged or is it optional?).

vi. As per claim 6: Lines 1-7: It is uncertain what is meant by "said mapping… is performed by passing I/O setup information… and sending data directly between said client partition requesting I/O access and said intelligent physical I/O hardware via said high performance memory channel" (i.e. Is this claiming a virtualization system or method steps? How is data sent directly? Does this mean it's sent without the intervention of the I/O partition? What about the monitor?). Line 4: It is uncertain what is meant by "a high performance memory channel" (i.e. What are its characteristics? What makes it high performance?). Lines 4-5: It is uncertain what is meant by "a client partition" (i.e. Does this refer to guest partition? Client partition?). Line 5: It is uncertain what is meant by "intelligent physical I/O hardware" (i.e. What are its characteristics? What makes it intelligent?).

vii. As per claim 7: Line 2: It is uncertain what is meant by "said I/O channel server" (i.e. There is insufficient antecedent basis for this term - does this refer to the memory channel of claim 1 or the high performance memory channel of claim 6?).

viii. As per claim 8: Line 2: it is uncertain what is meant by "redundant I/O partitions" (i.e. Does this mean the partitions are identical? Do they share the same connections to the same I/O devices? Different I/O devices?).

ix. As per claim 9: Lines 6-7: It is uncertain what is meant by "maps said physical I/O hardware to endpoints of an I/O channel server" (i.e. What are endpoints of an I/O channel server? Do these correspond to the guest partitions or how the guest partitions access the I/O?).

x. As per claim 10, it is dependent on claim 9 but does not overcome the deficiencies of claim 9; therefore, it is rejected for the same reasons.

xi. As per claim 11: Lines 4-5: It is uncertain what is meant by "from multiple partitions connected to said at least one I/O partition by respective memory channels" (i.e. Do the multiple partitions refer only to guest partitions or also to I/O partitions?).

xii. As per claim 12: Line 3: It is uncertain what is meant by "client partition relative addresses" (i.e. Does client partition refer to guest partition?). Line 4: It is uncertain what is meant by "client memory channel drivers" (i.e. Does this relate to the memory channel comprising memory shared between the I/O and guest partition?). Line 5: It is uncertain what is meant by "the client partition" (i.e. Does this refer to the guest partition? I/O partition?).

xiii. As per claim 13: line 2: It is uncertain what is meant by "a server of said I/O partition" (i.e. Does this refer to the I/O channel server of claim 9?).

xiv. As per claim 14: Line 4: It is uncertain what is meant by "a high performance memory channel" (i.e. What are its characteristics? What makes it high performance?); It is uncertain what is meant by "a client partition" (i.e. Does this refer to guest partition?). Line 5: It is uncertain what is meant by "intelligent physical I/O hardware" (i.e. What are its characteristics? What makes it intelligent?). Lines 5-7: It is uncertain what is meant by "sending data directly between said client partition requesting I/O access and said intelligent physical I/O hardware via said high performance memory channel" (i.e. How is data sent directly? Does this mean it's sent without the intervention of the I/O partition? What about the monitor?).

xv. As per claim 15, it is dependent on claim 14 but does not overcome the deficiencies of claim 14; therefore, it is rejected for the same reasons.

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

5. Claims 1- 8 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

6. As per claim 1, it recites an "a virtualization system"; however, it appears that the virtualization system would reasonably be interpreted by one of ordinary skill in the art as software, per se, failing to be tangibly embodied or include any recited hardware as part of the system. Software alone is directed to a non-statutory subject matter. Applicant is advised to amend the claims to include hardware (i.e. processor and memory) to overcome the §101 rejection.

7. As per claims 2-8, they are dependent on claim 1 but do not overcome the deficiencies of claim 1; therefore, they are rejected for the same reasons.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

8. Claims 1-8 are rejected under 35 U.S.C. 103(a) as being unpatentable over Arndt (US 2003/0204648) in view of Kjos *et al.* (US 2004/0064668; hereinafter Kjos).

9. As per claim 1, Arndt teaches a virtualization system for a host computer having at least one host processor and system resources including physical I/O hardware and memory *("sharing I/O facilities among logical partitions", abstract, lines 1-2; figure 1, data processing system 100 having processors 101-104 and I/O adapters 120-121; "Each PCI I/O adapter 120-121 provides an interface between data processing system 100 and I/O devices", paragraph 23, lines 10-12)*, the system comprising:

- virtualization software that divides said host computer into a plurality of virtual partitions including at least one user guest partition that provides a virtualization environment for at least one guest operating system and at least one input/output (I/O) partition that maps said physical I/O hardware to endpoints of an I/O channel server in said at least one I/O partition, said I/O channel server sharing the physical I/O hardware with at least one guest partition via a memory channel comprising memory shared between said at least one I/O partition and said at least one guest partition *("Partition management firmware 210 performs a number of functions and services for partitions… to create and enforce the partitioning of logically partitioned platform 200", paragraph 31, lines 1-4; figure 2, partition 203 running OS 202; "hosting partition is the logical partition owning the physical I/O facilities making them available to the hosted partition as virtual I/O facilities", paragraph 9, lines 15-17)*;
- a resource database for use in managing use of said at least one host processor and said system resources *("TCE tables within logically partitioned platforms are maintained by a platform firmware component known as the hypervisor, which*

insures that the data within the TCE directs the DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 4-8); and

- at least one monitor that maintains guest applications in said at least one guest partition within memory space *("hypervisor 210 performs a number of functions and services for partitions… to create and enforce the partitioning of logically partitioned platform 200", paragraph 31, lines 1-4; "Hosting partition 300 and hosted partition 320 are logical partitions or segments of memory", paragraph 40, lines 1-2)*; and

but is silent on an input/output (I/O) partition that maps said physical I/O hardware to endpoints of an I/O channel server in said at least one I/O partition, said I/O channel server sharing the physical I/O hardware with at least one guest partition via a memory channel comprising memory shared between said at least one I/O partition and said at least one guest partition; memory divided into most privileged system memory and less privileged user memory; virtualization software that operates in said less privileged user memory; at least one monitor that operates in said most privileged system memory; at least one guest partition within memory space specified in said resource database; and a context switch between said at least one monitor and said respective guest and I/O partitions for controlling multitask processing of software in said partitions on said at least one host processor.

It is obvious to one of ordinary skill in the art at the time the invention was made for an input/output (I/O) partition that maps said physical I/O hardware to endpoints of an I/O channel server in said at least one I/O partition, said I/O channel server sharing the physical I/O hardware with at least one guest partition via a memory channel

comprising memory shared between said at least one I/O partition and said at least one guest partition because Arndt teaches the hosting partition, or the I/O partition, maps I/O page buffers to translation control entries and shares memory with the hosted partition to share its physical I/O adapter and physical I/O with the shared partition *("Hosting partition 300 uses standard TCE table 310 to manage the I/O buffer pages in memory 305"; paragraph 33, lines 2-4; "TCE directs DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 7-8; "virtual I/O facilities of the hosted partition are realized by the hosting partition's physical I/O facilities, accessing the hosted partition's memory using the data stored in the standard translation control entry table on the hosting partition", paragraph 8, lines 18-22)*. One of ordinary skill in the art at the time the invention would have been motivated to do this to improve sharing of physical I/O facilities among logical partitions *(paragraph 7, lines 2-4)*.

Kjos teaches

- memory divided into most privileged system memory and less privileged user memory *("FIG. 13 illustrates the 79-bit virtual-memory address space reserved for the monitor… monitor further partitions its own virtual-memory space into… guest physical regions… representing the physical memory for each guest operating system", paragraph 62, lines 6-13; "The monitor runs at the most privileged privilege level… guest operating systems operate at less privileged privilege levels", paragraph 6, lines 18-21; it is obvious that the memory space for the monitor is most privileged and the memory space for guest operating systems is less privileged)*;

- virtualization software that operates in said less privileged user memory *("guest physical regions… representing the physical memory for each guest operating system", paragraph 62, lines 9-13; "guest operating systems operate at less privileged privilege levels", paragraph 6, lines 20-21)*;
- at least one monitor that operates in said most privileged system memory *("FIG. 13 illustrates the 79-bit virtual-memory address space reserved for the monitor… monitor further partitions its own virtual-memory space into… guest physical regions… representing the physical memory for each guest operating system", paragraph 62, lines 6-13; "The monitor runs at the most privileged privilege level… guest operating systems operate at less privileged privilege levels", paragraph 6, lines 18-21)*; and
- at least one guest partition within memory space specified in said resource database *("guest physical regions… representing the physical memory for each guest operating system", paragraph 62, lines 9-13; "guest operating systems operate at less privileged privilege levels", paragraph 6, lines 20-21)*.

The combination of Arndt in view of Kjos teaches a context switch between said at least one monitor and said respective guest and I/O partitions for controlling multitask processing of software in said partitions on said at least one host processor *(Kjos: "The monitor needs to maintain state information for each concurrently executing guest operating system in order to context switch between guest operating systems… monitor restores any other registers to the values they had at the point that the incoming guest operating system was last context-switched out", paragraph 72, lines 1-41. It is obvious*

that there is a context switch between the hypervisor, the hosted partition and the hosting partition, or the monitor, and respective guest and I/O partitions).

One of ordinary skill in the art at the time the invention was made would have been motivated the modify Arndt with the teachings of Kjos to provide "virtual address spaces to multiple guest operating systems concurrently running on a host computer using host-computer machine-architectural support for virtual-memory-address-translation caching and translation cache searching" *(paragraph 1, lines 3-7)*.

10. As per claim 2, Arndt in view of Kjos teaches the system of claim 1, and Arndt further teaches wherein upon receipt of a request to said I/O channel server from said at least one guest partition to access physical I/O hardware said I/O partition checks with partition descriptors stored in a monitor associated with said at least one guest partition to verify that the requested physical I/O hardware access is valid *("hosting partition virtual device driver requests that the hosting partition OS convert the RTCE opaque handle and I/O address range into a 'cross memory descriptor' which will be de-referenced in step 555", paragraph 47, lines 3-6)*.

11. As per claim 3, Arndt in view of Kjos teaches the system of claim 1, and Arndt further teaches wherein said mapping by said at least one I/O partition of said physical I/O hardware of said host computer to endpoints Of said I/O channel server in said I/O partition is performed by I/O partition software that multiplexes through shared common I/O physical hardware any I/O requests to said common I/O physical hardware from

multiple partitions connected to said I/O partition by respective memory channels *(figure 5, steps 535-570; "convert the RTCE opaque handle and I/O address range into a 'cross memory descriptor'… hosting partition's physical I/O device driver requests that the hosting partition's OS map the I/O page buffers represented by the cross memory descriptor to a range of I/O addresses assigned by the virtual I/O device driver… hosting partition's physical device driver notifies its I/O device to perform the requested operation", paragraph 47, lines 3-22)*.

12. As per claim 4, Arndt in view of Kjos teaches the system of claim 3, and Arndt further teaches wherein an I/O monitor associated with said I/O partition implements a system call interface between said I/O monitor and said I/O partition, said system call interface converting and validating client partition relative addresses, obtained as buffer parameters of requests sent through respective memory channels from client memory channel drivers, as valid hardware physical addresses of memory currently assigned to the client partition requesting access to said common I/O physical hardware *("virtual I/O device driver requests its operating system to map the I/O buffer pages associated with the I/O request via the RTCE table associated with the virtual I/O device to a range of I/O addresses assigned by the virtual I/O device driver... validates the I/O buffer page addresses", paragraph 46, lines 3-9; "hosting partition virtual device driver requests that the hosting partition OS convert the RTCE opaque handle and I/O address range into a 'cross memory descriptor' which will be de-referenced in step 555", paragraph 47, lines 3-6)*.

13. As per claim 5, Arndt in view of Kjos teaches the system of claim 4, and Arndt further teaches wherein messages between a server of said I/O partition and said respective guest partitions are queued by the guest partitions and de-queued by the I/O partition server and the partition relative physical addresses are converted by the I/O partition server to physical I/O hardware addresses with the aid of the I/O monitor, whereby data may be exchanged with hardware I/O adapters connected between said I/O monitor and said common I/O physical hardware *(figure 5, steps 510-580; paragraphs 46-47)*.

14. As per claim 6, Arndt in view of Kjos teaches the system of claim 1, and Arndt further teaches wherein said mapping by said at least one I/O partition of said physical I/O hardware of said host computer to endpoints of said I/O channel server in said I/O partition is performed by passing I/O setup information via said memory channel to said I/O channel server so as to set up a high performance memory channel between a client partition requesting I/O access and intelligent physical I/O hardware and sending data directly between said client partition requesting I/O access and said intelligent physical I/O hardware via said high performance memory channel *("Hosting partition 300 uses standard TCE table 310 to manage the I/O buffer pages in memory 305"; paragraph 33, lines 2-4; "TCE directs DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 7-8; "hosting partition's physical I/O device driver requests that the hosting partition's OS map the I/O page buffers represented by the*

cross memory descriptor to a range of I/O addresses assigned by the virtual I/O device driver… hosting partition's physical device driver notifies its I/O device to perform the requested operation", paragraph 47, lines 11-22).

15. As per claim 7, Arndt in view of Kjos teaches the system of claim 6, and Arndt further teaches wherein the guest partition requesting I/O access transfers data via said I/O memory channel to said intelligent physical I/O hardware using one of a user mode I/O or direct memory access data transfer operation *("I/O devices that uses Direct Memory Access (DMA) facilities", paragraph 4, lines 16-17; "TCE directs DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 7-8)*.

16. As per claim 8, Arndt in view of Kjos teaches the system of claim 1, and Arndt further teaches wherein the at least one I/O partition includes two redundant I/O partitions *("multiple instances of a single operating system", paragraph 19, line 4)*.

17. Claims 9-15 are rejected under 35 U.S.C. 103(a) as being unpatentable over Arndt.

18. As per claim 9, Arndt teaches a method of managing I/O operations of a plurality of operating system instances on a host computer having at least one host processor and system resources including physical I/O hardware *("sharing I/O facilities among logical partitions", abstract, lines 1-2; figure 1, data processing system 100 having*

processors 101-104 and I/O adapters 120-121; "Each PCI I/O adapter 120-121 provides an interface between data processing system 100 and I/O devices", paragraph 23, lines 10-12), the method comprising the steps of:

- dividing said host computer into a plurality of virtual partitions including at least one user guest partition that provides a virtualization environment for at least one guest operating system and at least one input/output (I/O) partition *("Partition management firmware 210 performs a number of functions and services for partitions… to create and enforce the partitioning of logically partitioned platform 200", paragraph 31, lines 1-4; figure 2, partition 203 running OS 202; "hosting partition is the logical partition owning the physical I/O facilities making them available to the hosted partition as virtual I/O facilities", paragraph 9, lines 15-17)*; and
- tracking allocation of said memory channel in a resource database *("TCE tables within logically partitioned platforms are maintained by a platform firmware component known as the hypervisor, which insures that the data within the TCE directs the DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 4-8; "hypervisor copies selected entries from the remote translation control entry table into selected entries in a standard translation control entry table maintained for a hosting partition per requests made by the hosting partition's Operating system", paragraph 8, lines 11-15)*

but is silent on an input/output (I/O) partition that maps said physical I/O hardware to endpoints of an I/O channel server in said at least one I/O partition; said I/O channel server sharing the physical I/O hardware with at least one guest partition via a memory

channel comprising memory shared between said at least one I/O partition and said at least one guest partition.

It is obvious to one of ordinary skill in the art at the time the invention was made for an input/output (I/O) partition that maps said physical I/O hardware to endpoints of an I/O channel server in said at least one I/O partition; said I/O channel server sharing the physical I/O hardware with at least one guest partition via a memory channel comprising memory shared between said at least one I/O partition and said at least one guest partition because Arndt teaches the hosting partition, or the I/O partition, maps I/O page buffers to translation control entries and shares memory with the hosted partition to share its physical I/O adapter and physical I/O with the shared partition *("Hosting partition 300 uses standard TCE table 310 to manage the I/O buffer pages in memory 305"; paragraph 33, lines 2-4; "TCE directs DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 7-8; "virtual I/O facilities of the hosted partition are realized by the hosting partition's physical I/O facilities, accessing the hosted partition's memory using the data stored in the standard translation control entry table on the hosting partition", paragraph 8, lines 18-22)*. One of ordinary skill in the art at the time the invention would have been motivated to do this to improve sharing of physical I/O facilities among logical partitions *(paragraph 7, lines 2-4)*.

19. As per claim 10, Arndt teaches the method of claim 9, further comprising the step of checking with partition descriptors stored in a monitor associated with said at least one guest partition to verify that a requested physical I/O hardware access is valid

("hosting partition virtual device driver requests that the hosting partition OS convert the RTCE opaque handle and I/O address range into a 'cross memory descriptor' which will be de-referenced in step 555", paragraph 47, lines 3-6; "virtual I/O device driver requests its operating system to map the I/O buffer pages associated with the I/O request via the RTCE table associated with the virtual I/O device to a range of I/O addresses assigned by the virtual I/O device driver… validates the I/O buffer page addresses", paragraph 46, lines 3-9).

20. As per claim 11, Arndt teaches the method of claim 9, wherein said mapping by said at least one I/O partition of said physical I/O hardware of said host computer to endpoints of said I/O channel server in said I/O partition includes the steps of multiplexing through shared common I/O physical hardware any I/O requests to said common I/O physical hardware from multiple partitions connected to said at least one I/O partition by respective memory channels *(figure 5, steps 535-570; "convert the RTCE opaque handle and I/O address range into a 'cross memory descriptor'… hosting partition's physical I/O device driver requests that the hosting partition's OS map the I/O page buffers represented by the cross memory descriptor to a range of I/O addresses assigned by the virtual I/O device driver… hosting partition's physical device driver notifies its I/O device to perform the requested operation", paragraph 47, lines 3-22).*

21. As per claim 12, Arndt teaches the method of claim 11, further comprising the steps of implementing a system call interface between an I/O monitor and said I/O

partition and said system call interface converting and validating client partition relative addresses, obtained as buffer parameters of requests sent through respective memory channels from client memory channel drivers, as valid hardware physical addresses of memory currently assigned to the client partition requesting access to said common I/O physical hardware *("virtual I/O device driver requests its operating system to map the I/O buffer pages associated with the I/O request via the RTCE table associated with the virtual I/O device to a range of I/O addresses assigned by the virtual I/O device driver... validates the I/O buffer page addresses", paragraph 46, lines 3-9; "hosting partition virtual device driver requests that the hosting partition OS convert the RTCE opaque handle and I/O address range into a 'cross memory descriptor' which will be de-referenced in step 555", paragraph 47, lines 3-6)*.

22. As per claim 13, Arndt teaches the method of claim 12, further comprising the steps of the guest partitions queuing messages between a server of said I/O partition and said respective guest partitions, the I/O partition server de-queuing the partition relative physical addresses are converting the partition relative physical addresses to physical I/O hardware addresses, and exchanging data with hardware I/O adapters connected between said I/O monitor and said common I/O physical hardware *(figure 5, steps 510-580; paragraphs 46-47)*.

23. As per claim 14, Arndt teaches the method of claim 9, wherein the step of mapping by said at least one I/O partition of said physical I/O hardware of said host

computer to endpoints of said I/O channel server in said I/O partition includes the steps of passing I/O setup information via said memory channel to said I/O channel server so as to set up a high performance memory channel between a client partition requesting I/O access and intelligent physical I/O hardware and sending data directly between said client partition requesting I/O access and said intelligent physical I/O hardware via said high performance memory channel *("Hosting partition 300 uses standard TCE table 310 to manage the I/O buffer pages in memory 305"; paragraph 33, lines 2-4; "TCE directs DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 7-8; "hosting partition's physical I/O device driver requests that the hosting partition's OS map the I/O page buffers represented by the cross memory descriptor to a range of I/O addresses assigned by the virtual I/O device driver… hosting partition's physical device driver notifies its I/O device to perform the requested operation", paragraph 47, lines 11-22)*.

24. As per claim 15, Arndt teaches the method of claim 14, wherein the step of sending data directly between said client partition and said intelligent physical UO hardware comprises the step of using one of a user mode I/O and a direct memory access data transfer operation *("I/O devices that uses Direct Memory Access (DMA) facilities", paragraph 4, lines 16-17; "TCE directs DMA accesses of I/O adapters to the proper target memory", paragraph 5, lines 7-8)*.

Conclusion

25. Any inquiry concerning this communication or earlier communications from the examiner should be directed to BRIAN CHEW whose telephone number is (571)270-5571. The examiner can normally be reached on Monday-Thursday, 8:00AM-5:00PM EST.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Meng-Ai An can be reached on (571)272-3756. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Meng-Ai An/
Supervisory Patent Examiner, Art Unit 2195

/B. C./
Examiner, Art Unit 2195